



US005935207A

United States Patent [19][11] **Patent Number:** **5,935,207****Logue et al.**[45] **Date of Patent:** **Aug. 10, 1999**

[54] **METHOD AND APPARATUS FOR PROVIDING REMOTE SITE ADMINISTRATORS WITH USER HITS ON MIRRORED WEB SITES**

[75] Inventors: **Jay D. Logue**, San Jose; **Lee S. Mighdoll**, San Francisco, both of Calif.

[73] Assignee: **WebTV Networks, Inc.**, Mountain View, Calif.

[21] Appl. No.: **08/827,643**

[22] Filed: **Apr. 9, 1997**

Related U.S. Application Data

[63] Continuation-in-part of application No. 08/656,924, Jun. 3, 1996.

[51] Int. Cl.⁶ **G06F 15/16**

[52] U.S. Cl. **709/219; 709/203; 711/118**

[58] Field of Search **395/200.41, 200.42, 395/200.43, 200.44, 200.45, 200.46, 200.47, 200.48, 200.49, 200.54, 200.59; 705/10, 26; 707/513; 711/118, 124; 709/203, 229, 217, 218, 219; 710/10**

[56] **References Cited****U.S. PATENT DOCUMENTS**

5,325,423	6/1994	Lewis	379/93.08
5,488,411	1/1996	Lewis	348/8
5,538,255	7/1996	Barker	463/41
5,558,339	9/1996	Perlman	463/42
5,564,001	10/1996	Lewis	345/302
5,586,257	12/1996	Perlman	463/42
5,586,260	12/1996	Hu	395/187.01
5,612,730	3/1997	Lewis	348/8
5,675,510	10/1997	Coffey et al.	395/200.54
5,708,780	1/1998	Levergood et al.	395/200.59
5,712,979	1/1998	Grabber et al.	395/200.11
5,737,619	4/1998	Judson	707/500

5,751,956	5/1998	Kirsch	395/200.33
5,754,774	5/1998	Bittinger et al.	395/200.33
5,774,670	6/1998	Montulli	395/200.57
5,787,470	7/1998	DeSimone et al.	711/124
5,802,292	9/1998	Mogul	395/200.33
5,864,852	1/1999	Luotonen	707/10

OTHER PUBLICATIONS

Ari Loutonen et al., World-Wide Web Proxies, <http://www.w3.org/History/1994/Proxies/Overview.html>, pp. 1-20, May 24, 1994.

Caching Algorithm of CERN httpd, <http://www.w3.org/History/1994/Caching/Overview.html>, pp. 1-4, May 1994.

Mark L. Van Name et al., Proxy servers will change the Web, PC Week, Ziff-Davis Publishing Co., full text, Mar. 1996.

Henrik Frystyk et al., Tutorial on Server Administration, <http://www.w3.org/People/Frystyk/HttpdTutorial/>, pp. 1-19, Oct. 1994.

(List continued on next page.)

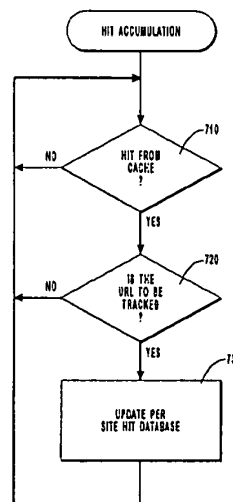
Primary Examiner—Zarni Maung

Assistant Examiner—Patrice L. Winder

Attorney, Agent, or Firm—Workman, Nydegger, Seeley

[57] **ABSTRACT**

A method and apparatus for providing mirrored site administrators with the number of hits from a proxy's document cache and for dispatching document requests in a proxy to more efficiently allocate the document cache space within the proxy are provided. A proxy includes a document cache storing recently requested documents. The proxy is coupled to a client and to a remote server. The proxy maintains information regarding requests from the client that are serviced from the proxy's document cache such as the Uniform Resource Locator (URL) of the requested document and the number of cached responses. This information is provided by the proxy to a remote site administrator. In this manner, remote site administrators can more accurately track total hits.

16 Claims, 10 Drawing Sheets

OTHER PUBLICATIONS

Martin Gleeson, pwebstats: A Perl Web Statistics Generator, <http://www.unimelb.edu.au/pwebstats/>, pp. 1–17, Oct. 1995.
Proxies, <http://www.w3.org/Daemon/User/Proxies/Proxies.html>, pp. 1–3, author unknown, Jul. 1995.
Ari Loutonen et al., CERN httpd PreRelease Notes, CERN, http://apollo.mchime-u.ac.jp/ReleaseNotes_3.Opre.html, pp. 1–3, Apr. 1994.
Ari Loutonen, Logging Control In CERN httpd, CERN, <http://www.uni-karlsruhe.de/Betrieb/www...aemon/User/Config/Logging.html#AccessLog/>, pp. 1–4, Dec. 1994.
Azer Bestavros et al., Application-Level Document Caching in the Internet, IEEE, pp. 166–173, Jun. 1995.

Clinton L. Jeffrey et al., Proxy-Sharing Proxy Servers, IEEE, pp. 116–119, Mar. 1996.

Jeffrey Mogul, DECWRL, Paul J. Leach, Microsoft, Simple Hit-Metering and Usage-Limiting for HTTP, HTTP Working Group, Internet-Draft, Mar. 19, 1997, pp. 1–37.

Administrator's Guide, Netscape Proxy Server, ©Netscape Communications Corp. 1995, 1996, pp. 19–20.

Chapman et al., Building Internal Firewalls Chapter 7: Proxy Systems O'Reilly & Associates, Inc., Sep. 1995, pp. 189–205.

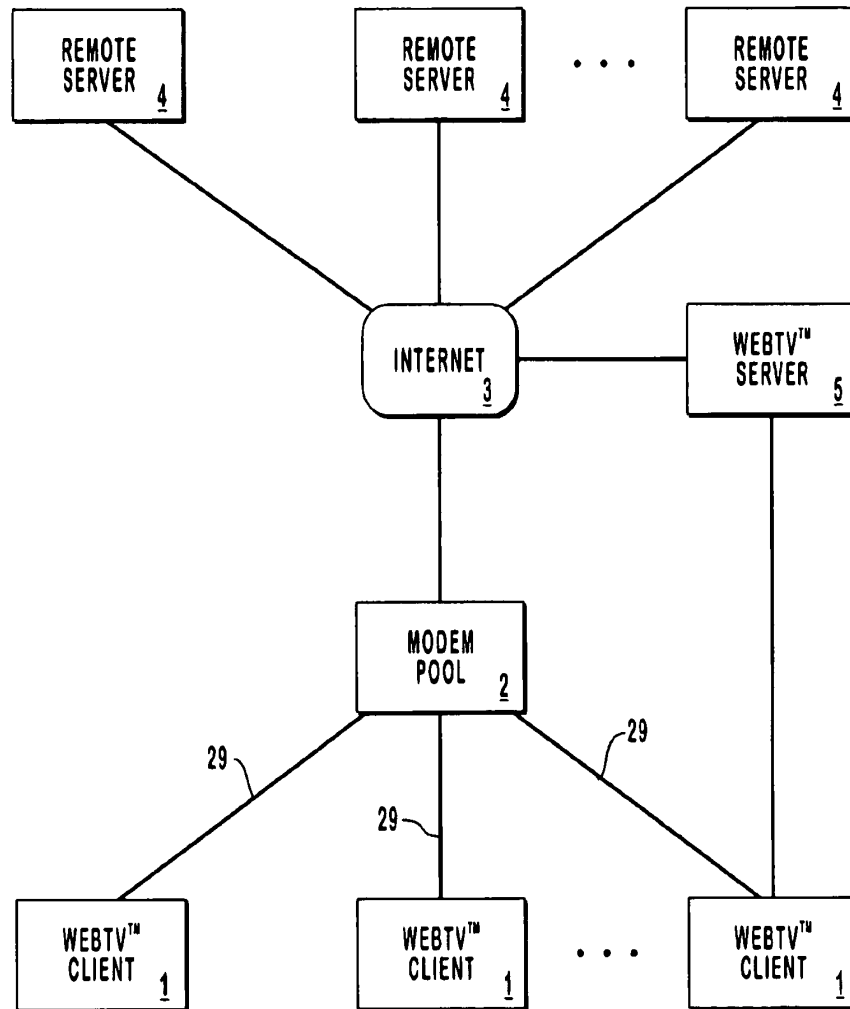


FIG. 1

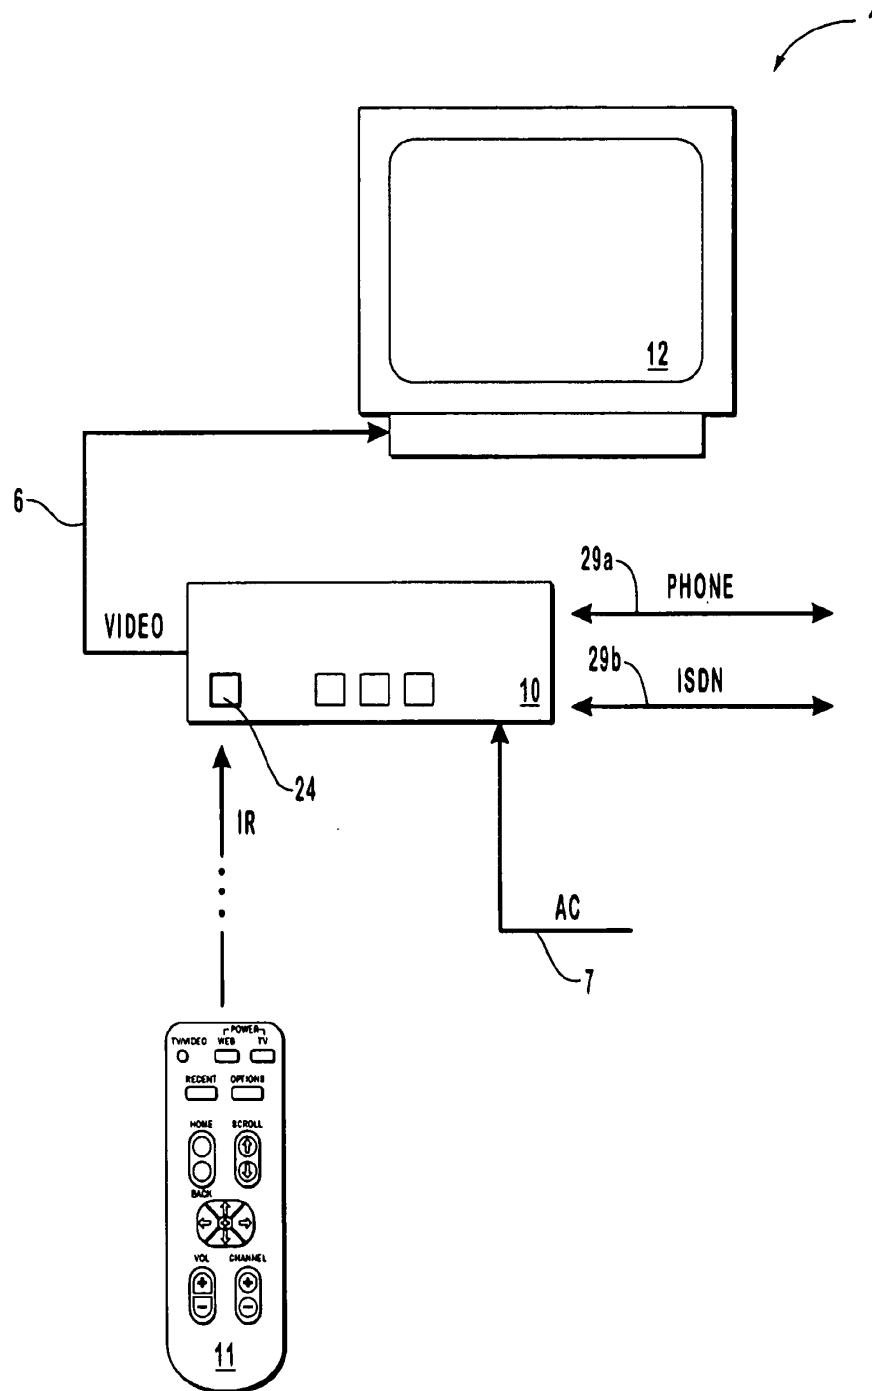


FIG. 2

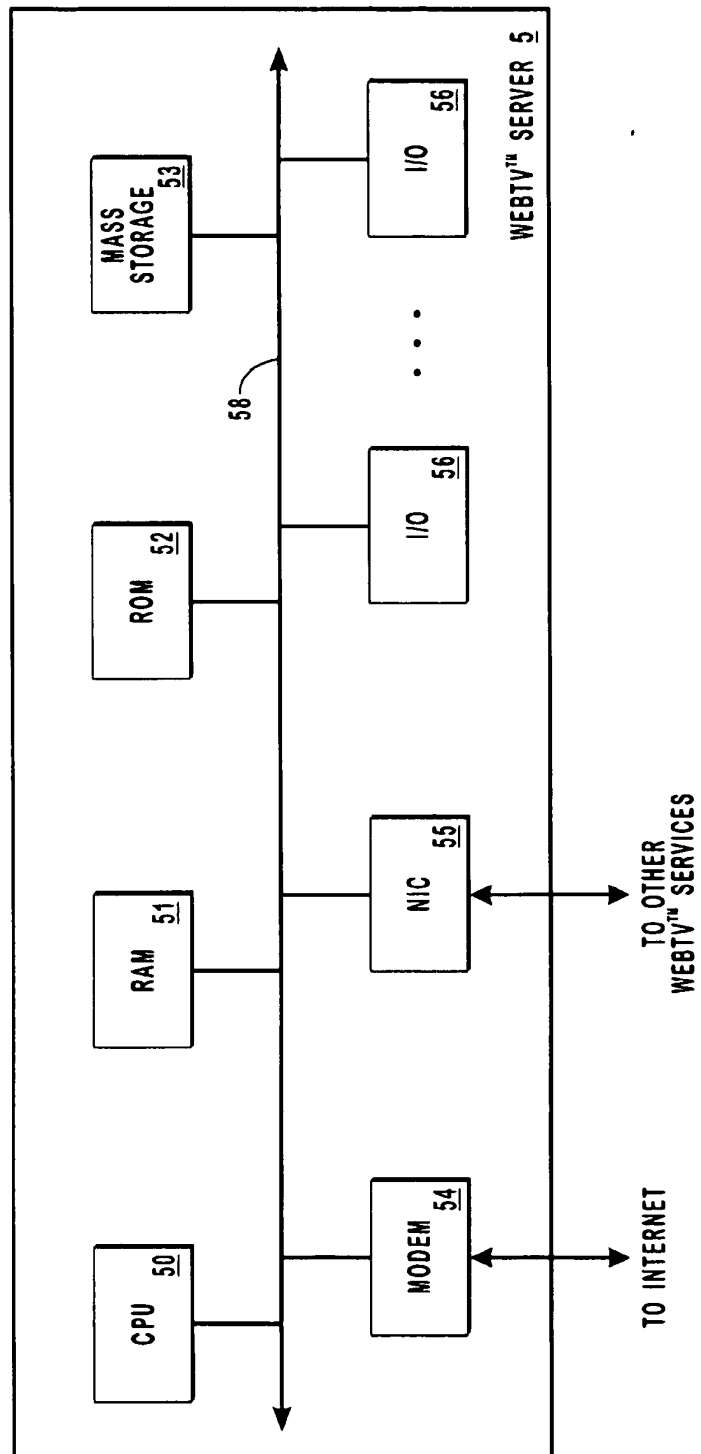


FIG. 3

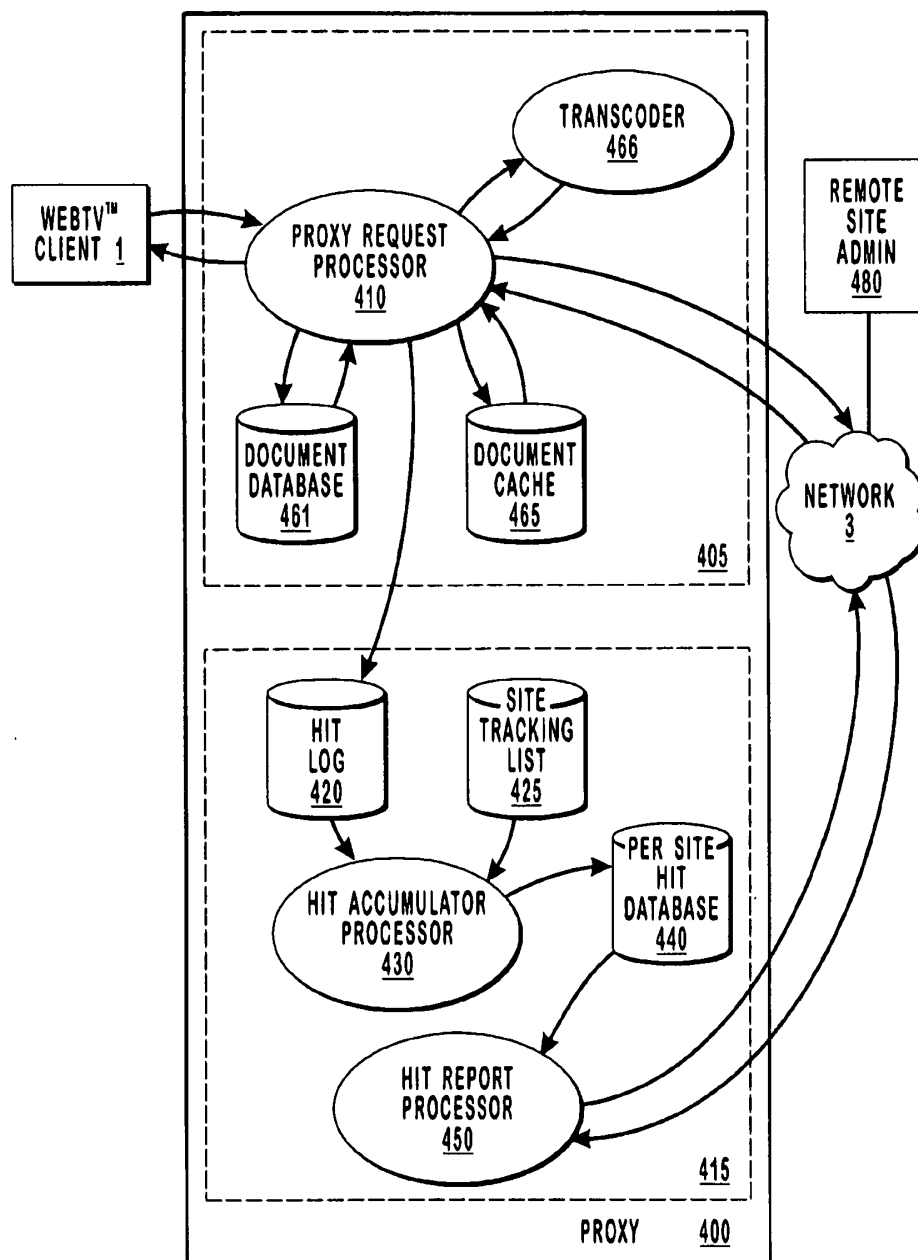


FIG. 4

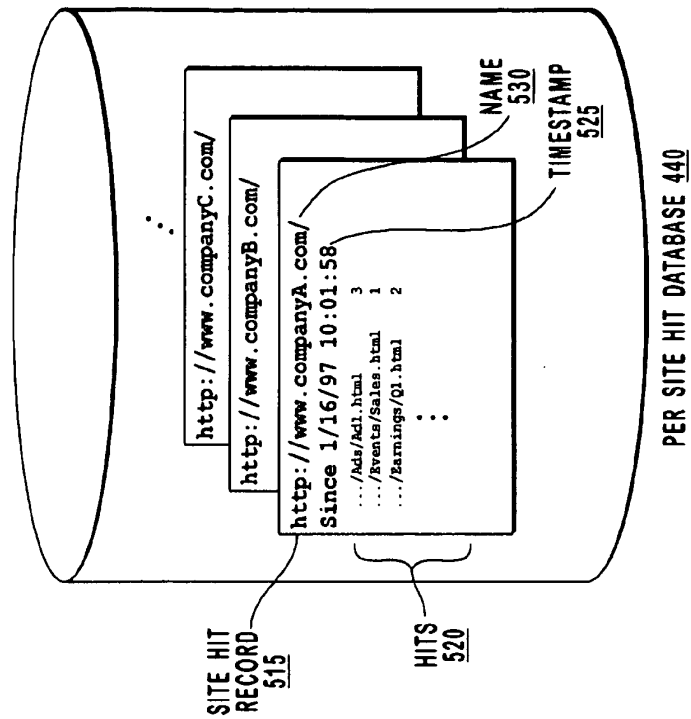


FIG. 5B

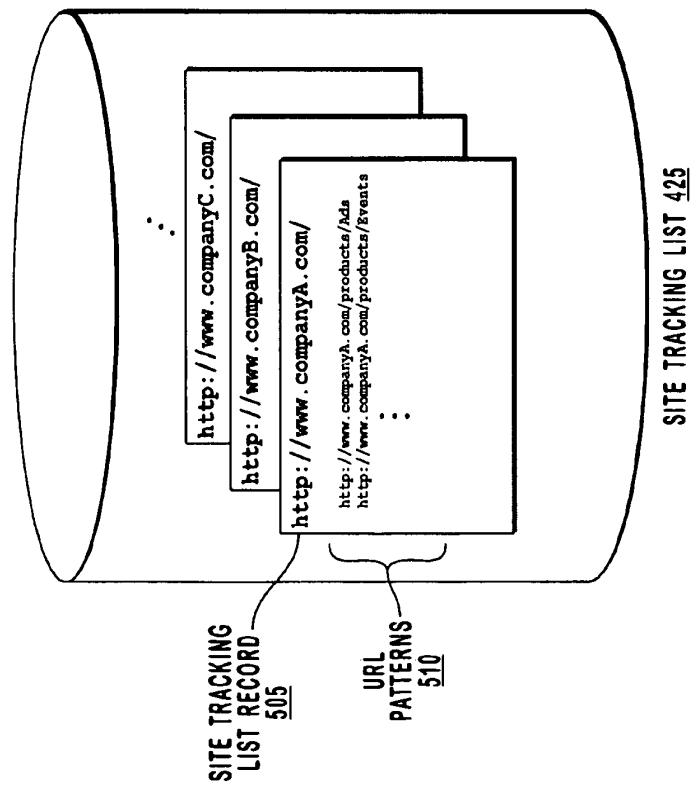


FIG. 5A

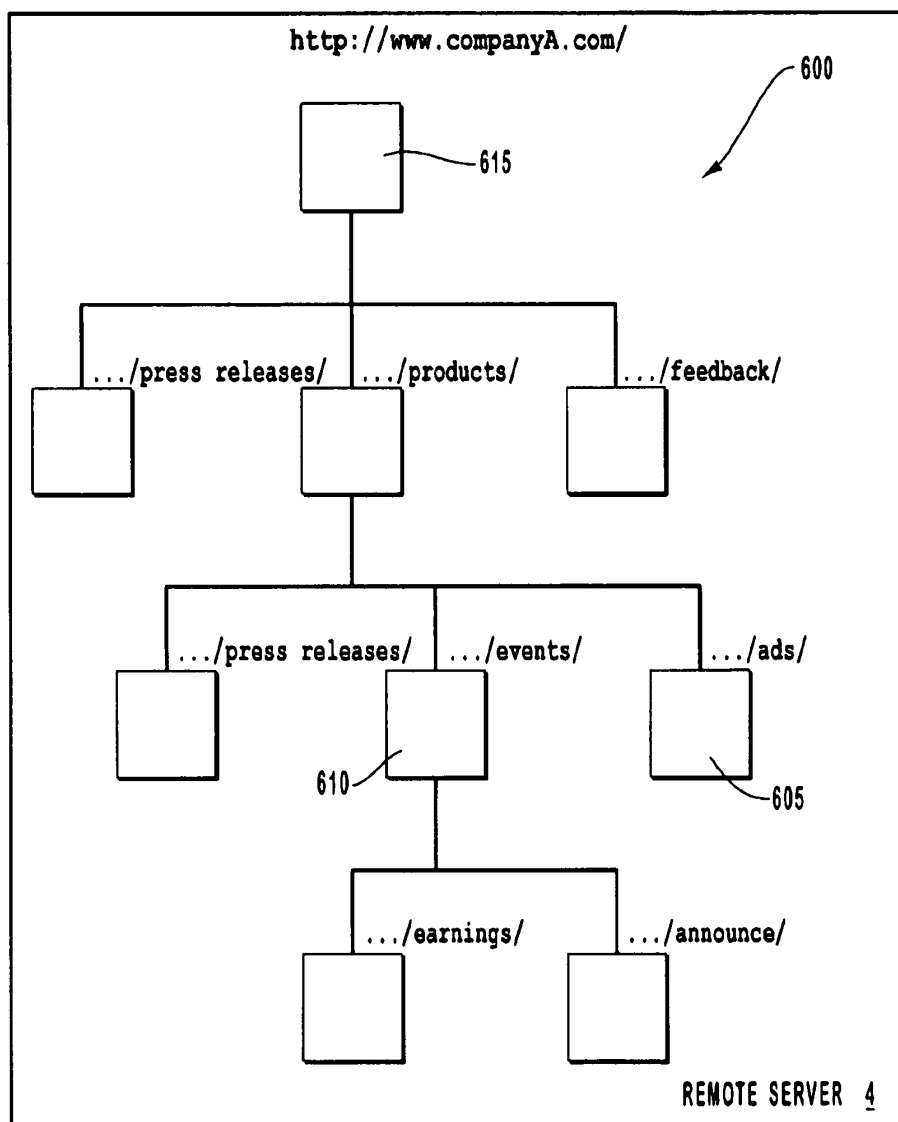


FIG. 6

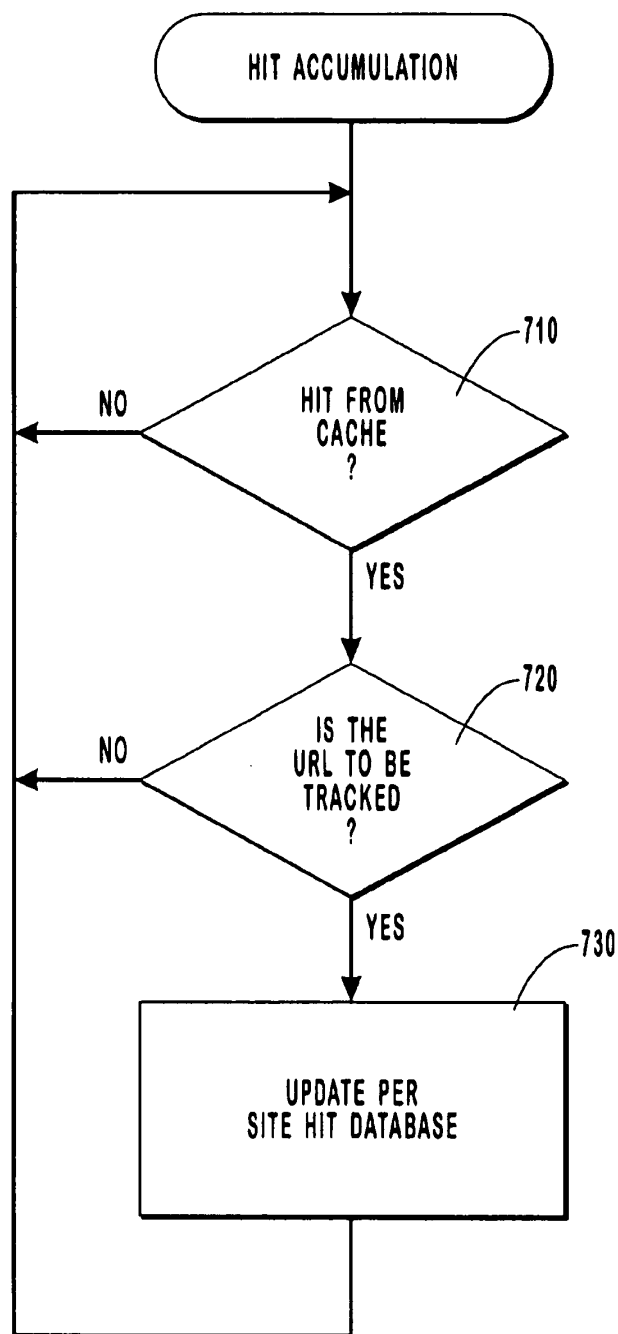


FIG. 7

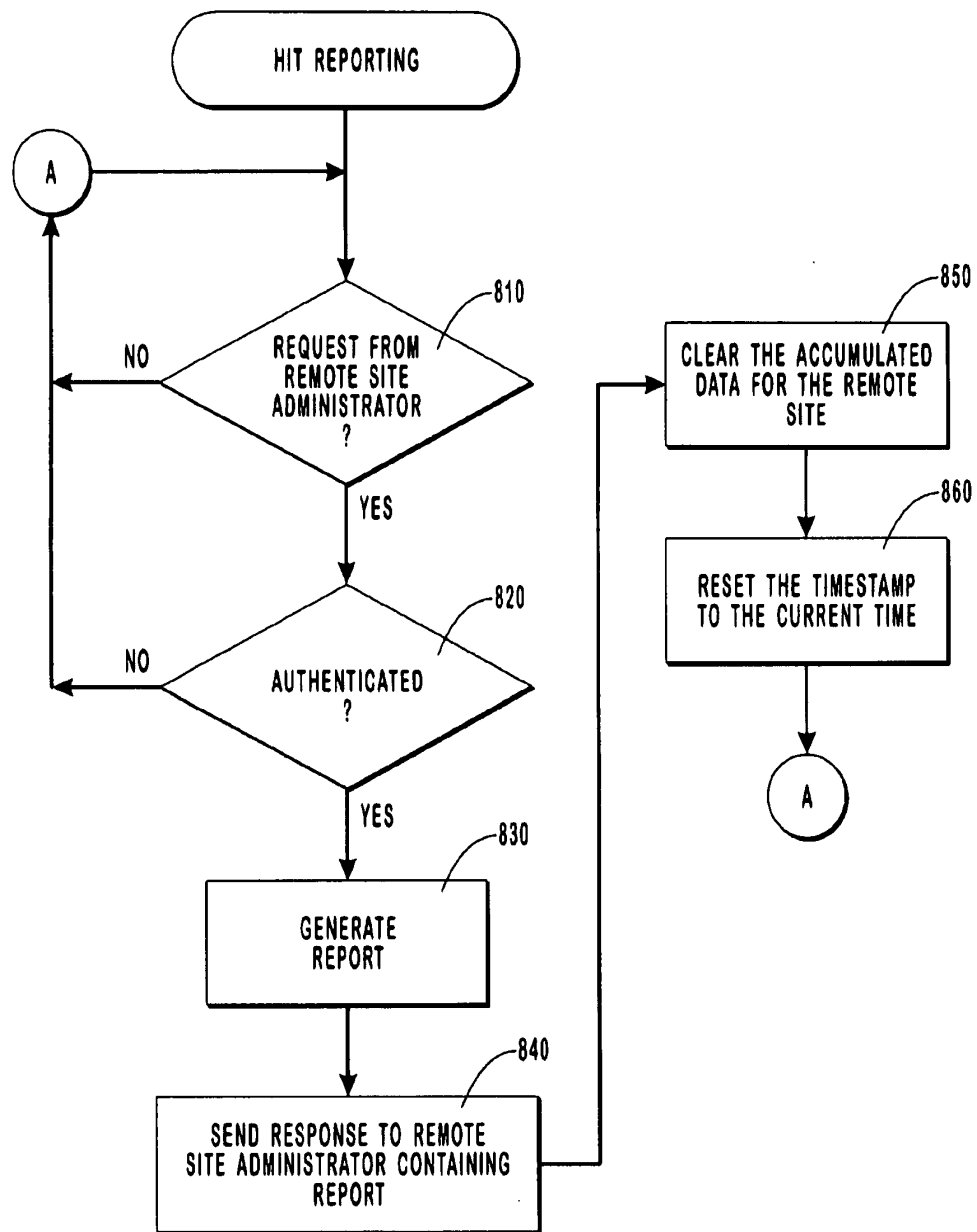


FIG. 8

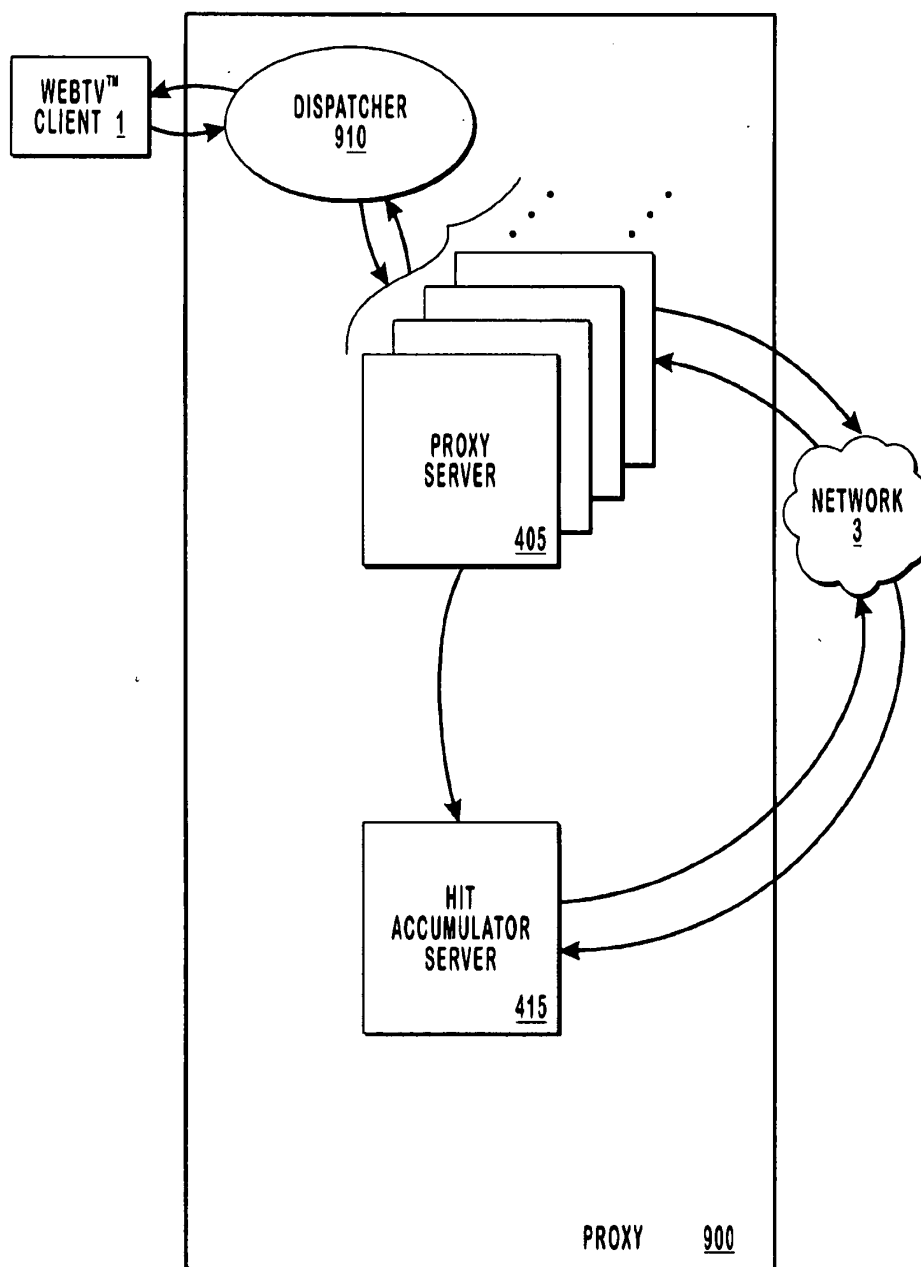


FIG. 9

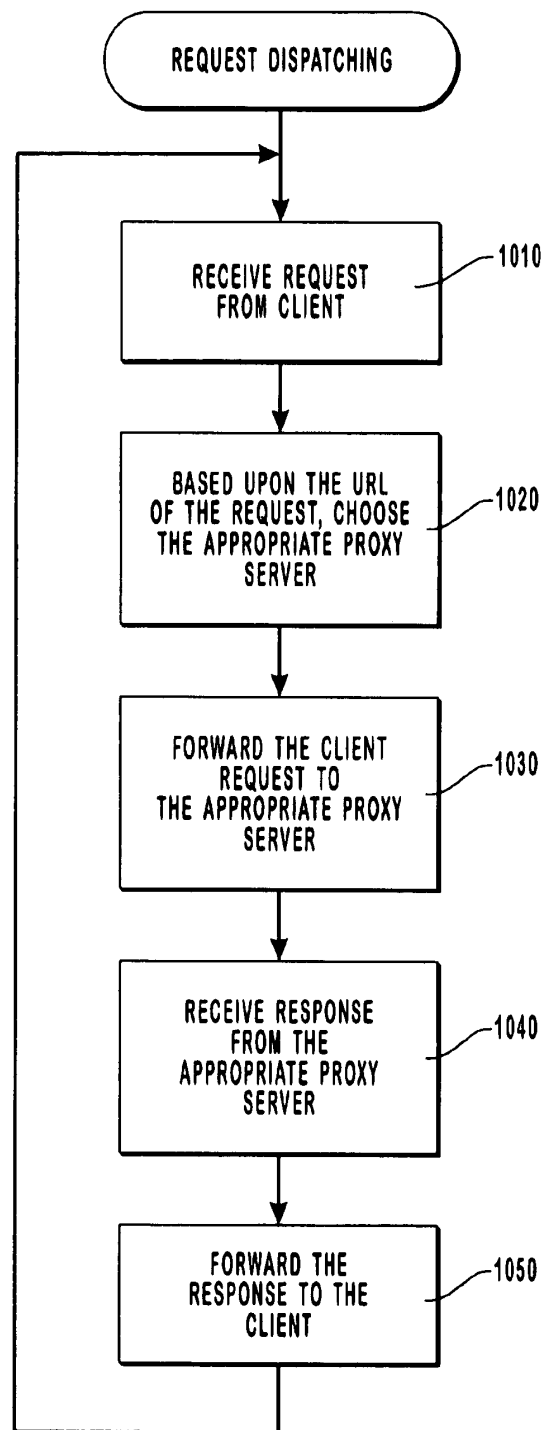


FIG. 10

METHOD AND APPARATUS FOR PROVIDING REMOTE SITE ADMINISTRATORS WITH USER HITS ON MIRRORED WEB SITES

CROSS-REFERENCES TO RELATED APPLICATIONS

The present application is a continuation-in-part of co-pending U.S. patent application entitled, "Method and Apparatus for Providing Proxying and Transcoding of Documents in a Distributed Network," having application Ser. No. 08/656,924, and filed on Jun. 3, 1996.

FIELD OF THE INVENTION

The invention relates generally to the field of client-server computer networking. More specifically, the invention relates to mirroring of Web sites, notifying mirrored site administrators of hits, and allocation of the Web's content among mirroring servers based upon the Uniform Resource Locator (URL).

BACKGROUND OF THE INVENTION

World Wide Web (Web) documents are commonly written in HTML (Hypertext Mark-up Language). HTML documents typically reside on Web servers and are requested by Web clients. Often delays can be introduced during Web browsing by heavy communications traffic on the Internet or slow response of a remote site, for example. Providing one or more servers for mirroring Web sites located on remote servers is one means of reducing delays involved with browsing the Web. These mirroring servers, typically referred to collectively as a "proxy" or individually as "proxy servers," store frequently accessed Web sites in a local cache, thereby eliminating recurrent retrievals of commonly accessed documents. Thus, when a request for a particular Web page is received from a client, the proxy server associated with the particular client looks first to its local cache to service the request rather than the remote site upon which the Web page resides. If the requested document is found locally, the request can be serviced by the proxy server and a subsequent request to the remote server for the document can be avoided. Therefore, only when a valid copy of the requested document is not in the proxy's local cache would the remote server need to be accessed. In this manner, exposure to heavy communications traffic on the Internet and slow response of remote serves can be reduced.

While this mirroring approach is beneficial to end-users, it makes hit tracking for remote site administrators difficult. A hit is a request for a Web page, typically initiated by a user selecting a hypertext link for the Web page. The mirroring approach discussed above disrupts a remote server's ability to track the total number of requests for a given Web page because, as discussed above, some of the requests are intercepted and serviced by proxy servers. It is desirable to have an accurate count of requests for a given Web page or group of pages to track the relative popularity of a page, for example or to provide feedback to advertisers whose advertisements appear on the page. Therefore, what is needed is a mechanism for tracking user hits by the proxy and a mechanism for notifying mirrored sites, thereby allowing remote site administrators to accurately track total hits (i.e., those requests serviced from a proxy's local cache and the requests serviced by the remote server).

Another problem with the current mirroring approach is the inefficient allocation of the proxy's cache space.

Currently, each client is assigned to one or more proxy servers. Therefore, the documents most recently requested by each active client will reside in the corresponding proxy server's cache. Assuming one or more clients assigned to different proxy servers have requested the same document recently, the same document might be cached in several of the proxy servers, thereby reducing the cache storage space for other frequently requested documents. Further, one or more extremely popular documents might potentially be cached in each proxy server. While redundancy of information is useful for fault tolerance, organized redundancy would be preferable. Given the foregoing, what is needed is a means of more efficiently allocating cache space within a proxy. Specifically, it would be desirable to allocate mutually exclusive portions of the Web's content to particular proxy servers.

SUMMARY OF THE INVENTION

A method and apparatus are described for providing mirrored site administrators with the number of hits from a proxy's document cache and for dispatching document requests in a proxy to more efficiently allocate the document cache space within the proxy. A proxy includes a document cache storing recently requested documents. The proxy is coupled to a client and to a remote server. The proxy maintains information regarding requests from the client that are serviced from the proxy's document cache. This information is provided by the proxy to a remote site administrator. In this manner, remote site administrators can more accurately track total hits (i.e., those requests serviced from a proxy's document cache plus the requests serviced by the remote server itself).

According to another aspect of the present invention, a proxy implements a dispatching scheme for client requests that results in a more efficient allocation of the proxy's document cache space. The proxy receives a document request from a client. A Uniform Resource Locator (URL) is included in the document request. The proxy forwards the request to one of a plurality of proxy servers based upon the URL.

According to another aspect of the present invention, the proxy performs a hash function on the URL that maps the URL to exactly one of the plurality of proxy servers. Advantageously, in this manner, mutually exclusive portions of the Web's content can be allocated to particular proxy servers.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention is illustrated by way of example, and not by way of limitation, in the figures of the accompanying drawings and in which like reference numerals refer to similar elements and in which:

FIG. 1 is a block diagram illustrating several clients connected to a proxy server in a network.

FIG. 2 is a diagram illustrating a client according to one embodiment of the present invention.

FIG. 3 is a block diagram of a server according to one embodiment of the present invention.

FIG. 4 is a data flow diagram illustrating the interaction of proxy components according to one embodiment of the present invention.

FIG. 5A is a depiction of an exemplary site tracking list according to one embodiment of the present invention.

FIG. 5B is a depiction of an exemplary per site hit database according to one embodiment of the present invention.

FIG. 6 is a logical view of an exemplary directory structure of a remote server.

FIG. 7 is a flow diagram illustrating a method of performing hit accumulation according to one embodiment of the present invention.

FIG. 8 is a flow diagram illustrating a method of hit reporting according to one embodiment of the present invention.

FIG. 9 is a data flow diagram illustrating the interaction of proxy components according to another embodiment of the present invention.

FIG. 10 is a flow diagram illustrating a method of dispatching requests to segregate the storage of documents according to one embodiment of the present invention.

DETAILED DESCRIPTION

A method and apparatus are described for providing mirrored site administrators with the number of hits from a proxy's document cache and for maintaining a more efficient document caching scheme in a client-server computer network. In the following description, for purposes of explanation, numerous specific details are set forth in order to provide a thorough understanding of the present invention. It will be evident, however, to one skilled in the art that the present invention may be practiced without these specific details. Further, in other instances, well-known structures and devices are shown in block diagram.

The present invention includes various steps, which will be described below. The steps can be embodied in machine-executable instructions, which can be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps of the present invention might be performed by specific hardware components that contain hardwired logic for performing the steps, or by any combination of programmed computer components and custom hardware components.

While embodiments of the present invention will be described with respect to HTML documents, the method and apparatus described herein are equally applicable to other types of documents such as text files, images (e.g., JPEG and GIF), audio files (e.g., .WAV, .AU, and .AIFF), video files (e.g., .MOV, and .AVI), and other document types commonly found on the Web.

System Overview

The present invention may be included in a system, known as WebTV™, for providing a user with access to the Internet. A user of a WebTV™ client generally accesses a WebTV™ server via a direct-dial telephone (POTS, for "plain old telephone service"), ISDN (Integrated Services Digital Network), or other similar connection, in order to browse the Web, send and receive electronic mail (e-mail), and use various other WebTV™ network services. The WebTV™ network services are provided by WebTV™ servers using software residing within the WebTV™ servers in conjunction with software residing within a WebTV™ client.

FIG. 1 illustrates a basic configuration of the WebTV™ network according to one embodiment. A number of WebTV™ clients 1 are coupled to a modem pool 2 via direct-dial, bi-directional data connections 29, which may be telephone (POTS, i.e., "plain old telephone service"), ISDN (Integrated Services Digital Network), or any other similar type of connection. The modem pool 2 is coupled typically through a router, such as that conventionally known in the

art, to a number of remote servers 4 via a conventional network infrastructure 3, such as the Internet. The WebTV™ system also includes a WebTV™ server 5, which specifically supports the WebTV™ clients 1. The WebTV™ clients 1 each have a connection to the WebTV™ server 5 either directly or through the modem pool 2 and the Internet 3. Note that the modem pool 2 is a conventional modem pool, such as those found today throughout the world providing access to the Internet and private networks.

Note that in this description, in order to facilitate explanation the WebTV™ server 5 is generally discussed as if it were a single device, and functions provided by the WebTV™ services are generally discussed as being performed by such single device. However, the WebTV™ server 5 may actually comprise multiple physical and logical devices connected in a distributed architecture, and the various functions discussed below which are provided by the WebTV™ services may actually be distributed among multiple WebTV™ server devices.

An Exemplary Client System

FIG. 2 illustrates a WebTV™ client 1. The WebTV™ client 1 includes an electronics unit 10 (hereinafter referred to as "the WebTV™ box 10"), an ordinary television set 12, and a remote control 11. In an alternative embodiment of the present invention, the WebTV™ box 10 is built into the television set 12 as an integral unit. The WebTV™ box 10 includes hardware and software for providing the user with a graphical user interface, by which the user can access the WebTV™ network services, browse the Web, send e-mail, and otherwise access the Internet.

The WebTV™ client 1 uses the television set 12 as a display device. The WebTV™ box 10 is coupled to the television set 12 by a video link 6. The video link 6 is an RF (radio frequency), S-video, composite video, or other equivalent form of video link. In the preferred embodiment, the client 1 includes both a standard modem and an ISDN modem, such that the communication link 29 between the WebTV™ box 10 and the server 5 can be either a telephone (POTS) connection 29a or an ISDN connection 29b. The WebTV™ box 10 receives power through a power line 7.

Remote control 11 is operated by the user in order to control the WebTV™ client 1 in browsing the Web, sending e-mail, and performing other Internet-related functions. The WebTV™ box 10 receives commands from remote control 11 via an infrared (IR) communication link. In alternative embodiments, the link between the remote control 11 and the WebTV™ box 10 may be RF or any equivalent mode of transmission.

An Exemplary Server System

The WebTV™ server 5 generally includes one or more computer systems generally having the architecture illustrated in FIG. 3. It should be noted that the illustrated architecture is only exemplary; the present invention is not constrained to this particular architecture. The illustrated architecture includes a central processing unit (CPU) 50, random access memory (RAM) 51, read-only memory (ROM) 52, a mass storage device 53, a modem 54, a network interface card (NIC) 55, and various other input/output (I/O) devices 56. Mass storage device 53 includes a magnetic, optical, or other equivalent storage medium. I/O devices 56 may include any or all of devices such as a display monitor, keyboard, cursor control device, etc. Modem 54 is used to communicate data to and from remote servers 4 via the Internet.

As noted above, the WebTV™ server 5 may actually comprise multiple physical and logical devices connected in a distributed architecture. Accordingly, NIC 55 is used to provide data communication with other devices that are part of the WebTV™ services. Modem 54 may also be used to communicate with other devices that are part of the WebTV™ services and which are not located in close geographic proximity to the illustrated device.

An Exemplary Proxy

FIG. 4 illustrates the caching and hit accumulation features of the WebTV™ proxy 400 according to one embodiment of the present invention. In this embodiment, one or more WebTV™ servers 5 may act as a proxy 400 in providing the WebTV™ client 1 with access to the Web and other WebTV™ services. More specifically, WebTV™ server 5 functions as a "caching proxy." In this example, proxy 400 includes a proxy server 405 and a hit accumulator server 415. Client requests that are serviced from the proxy server's local document cache 465 are communicated to the hit accumulator server 415. As will be described below, the hit accumulator server 415 maintains and organizes the data so as to provide hit tracking information to remote site administrators such as remote site administrator 480. Remote site administrator 480 may include entities such as persons authorized to gather statistical data for the remote site, persons authorized to manage and maintain the remote site the site, the remote site itself, or an automated computer system or other device configured to receive statistical data for the remote site.

In this embodiment, the proxy server 405, includes a proxy request processor 410, a document cache 465, a document database 461, and a transcoder 466. The proxy request processor 410 receives requests from the WebTV™ client 1 and sends responses to the WebTV™ client 1. The proxy request processor 410 maintains the document database 461, the document cache 465, and further determines when transcoding will be performed. The document cache 465 is used for temporary storage of Web documents such as images, text files, audio files, video files and other information which is used frequently by either WebTV™ client 1 or the proxy server 405.

When a document request is received from a client, the proxy request processor 410 determines whether to service the request from the document cache 465 by performing a search of the document cache 465. If the document is found locally, then the document may be retrieved from the document cache 465 and transferred to the client with the response. However, if the requested document is not found, then the proxy request processor 410 requests the document from the appropriate site and upon receipt the proxy request processor 410 provides the document to the client with the response. Further, the proxy request processor 410 anticipates subsequent requests by storing the document in the document cache 465.

When a document is retrieved by the proxy server 405 from a remote server 4, for example, detailed information on this document may be stored in the document database 461. The stored information may subsequently be used by the proxy server 405 to speed up processing and downloading of that document in response to future requests for that document. In addition, the transcoding functions and various other functions of the WebTV™ service may be facilitated by making use of information stored in the document database 461. For example, the document database 461 may include certain historical and diagnostic information for Web pages that have been accessed by a WebTV™ client 1.

Document transcoder 466 is used to automatically revise the code of Web documents retrieved from the remote servers 4, for purposes such as: (1) correcting bugs in documents; (2) correcting undesirable effects which occur when a document is displayed by the client 1; (3) improving the efficiency of transmission of documents from the server 5 to the client 1; (4) matching hardware decompression technology within the client 1; (5) resizing images to fit on the television set 12; (6) converting documents into other formats to provide compatibility; (7) reducing latency experienced by a client 1 when displaying a Web page with in-line images (images displayed in text); and, (8) altering documents to fit into smaller memory spaces.

In one embodiment, hit accumulator server 415 may act as a Web server providing a Hypertext Transport Protocol (HTTP) interface by which remote site administrators can access the accumulated hits for their sites by way of a Web browser. The hit accumulator server 415 may include a hit log 420, a hit accumulator processor 430, a site tracking list 425, a hit report processor 450, and a per site hit database 440. One method of communicating hits from a given proxy server to the hit accumulator server 415 is through a common storage device such as hit log 420. This and other methods of communicating hits will be described below. Regardless of how hits are communicated to the hit accumulator server 415, a process such as the hit accumulator processor 430 is desirable to verify the hits against a list of locations that are to be monitored. Such a list of locations may be stored in the site tracking list 425, for example. A location, in this context, refers to the location of a document. The location may be represented by a URL, a directory path, or other mechanisms for uniquely identifying a particular document. Hits that are validated by the hit accumulator processor 430 are recorded in the per site hit database 440. Thus, the per site hit database 440 will have a current count of the hits for each location listed in the site tracking list 425. In this embodiment, the hit report processor 450 may receive requests from remote site administrators such as remote site administrator 480 for hit reports. The hit reports can be extracted from the per site hit database 440 and transmitted to the requester in an HTML report, for example.

While in this embodiment the proxy server 405 and the hit accumulation server 415 have been shown as separate servers, the functionality of both could be combined into one WebTV™ server 5. Additionally, the proxy 400 might be expanded to include more than one proxy server 405. When expanding the proxy 400 to include more than one proxy server 405, only one hit accumulation server 415 need be employed.

In alternative embodiments, hits may be communicated by a proxy server 405 to the accumulation server 415 by way of a network connection such as permanent connection through which events may be sent. Also, message passing may be employed whereby the proxy server 405 sends a message such as a datagram to the hit accumulator 415 to notify it of a document cache hit. It is appreciated that many other means of communicating information between servers are possible.

An Exemplary Site Tracking List

FIG. 5A illustrates an exemplary site tracking list according to one embodiment of the present invention. This illustration depicts a site tracking list 435 including site tracking list records 505 for three remote sites: (1) <http://www.companyA.com/>; (2) <http://www.companyB.com/>; and (3) <http://www.companyC.com/>. In this embodiment,

each site tracking list record 505 may include a list of one or more URL patterns 510.

The list of URL patterns 510 may be a list of strings identifying the initial portions (e.g., prefixes) of URLs to be tracked. In this example, the proxy 400 tracks hits for documents identified by URLs with a prefix that matches any of the URL patterns 510 specified in one of the site tracking list records 505. The hits may then be logged to a record in the per site hit database 440 corresponding to the site tracking list record 505 which contained the matching URL pattern. This form of URL pattern is useful for tracking hits for a particular grouping of Web pages beginning with the same initial sequences of characters. For example, the URLs for the Web pages of Company A might all begin with "http://www.company_A.com/." Additionally, the Web pages associated with products produced by Company A might all begin with the sequence "http://www.company_A.com/product/." Furthermore, pages related to a particular product, might all begin with the URL prefix "http://www.company_A.com/product/<product_name>/" where <product_name> identifies the particular product. To track the hits for pages relating to Company A's Gizmo product line, therefore, the following URL pattern may be used: "http://www.company_A.com/product/Gizmo/." Similarly, to track the hits for all of Company A's products the following URL pattern may be used: "http://www.company_A.com/product/."

URL patterns are not limited to prefixes, other forms of URL patterns may be used such as patterns including wild card or other special characters, or patterns in the form of standard regular expressions.

An Exemplary Per Site Hit Database

FIG. 5B illustrates an exemplary per site hit database according to one embodiment of the present invention. Based upon the information provided in the site tracking list 425 of FIG. 5A, an exemplary per site hit database might be represented as per site hit database 440. In this example, the per site hit database 440 includes three sit hit records 515 corresponding to remote sites for CompanyA, CompanyB and CompanyC.

In this embodiment, each site hit record 515 includes a timestamp 525. The timestamp 525 may indicate the time from which the hits have been accumulated. In this example, therefore, there have been six hits to the monitored URLs since Jan. 16, 1997 at 10:01:58. Those of skill in the art will appreciate the timestamp 525 may represent the period of accumulation in other ways such as elapsed time since the last hit report was generated.

Site hit records 515 also include a remote site name 530. The remote site names 530 from front to back correspond to CompanyA, CompanyB, and CompanyC. Site hit record 515 further includes a list of hits 520. In this embodiment, the list of hits 520 includes the URLs of the documents that were requested and subsequently serviced from the proxy's local cache (e.g., document cache 465) since the time indicated by the timestamp 525. According to the site hit record 515 for CompanyA, the ad1.html Web page has been requested and serviced from the proxy's local cache three times. Similarly, the sales.html and Q1.html Web pages have been provided from the proxy's cache once and twice, respectively. Based upon the accumulated hit information in a particular site hit record 515, a detailed hit report may be provided to the corresponding remote site administrator. Hit accumulation will be discussed further below.

FIG. 6 is a logical view of an exemplary directory structure 600 that may exist on a remote server 4. This exemplary directory structure 600 illustrates the need for a flexible method of tracking the number of hits. Web pages might reside in any or all of the directories shown. In this example, the URL patterns within a site tracking list record 505 may identify a particular directory or directories in the hierarchy depicted.

The remote site administrator for CompanyA may want to know the number of hits in an Ads subdirectory 605 and an Events subdirectory 610. This may be due to the fact that advertising banners are shown on Web pages in these directories and the advertisers may want feedback on how many Web viewers are seeing their ads. Alternatively, the company may have its own business reasons for analyzing statistics in certain areas of their Web site. Regardless, it is apparent that simply tracking all hits for a root directory 615 on the company's server is insufficient. For example, hits would be tracked for directories in which the remote site administrator had no interest. A list of URL patterns is used to accommodate the flexibility desired. The following URL patterns may be stored in the site tracking list 425 for CompanyA to track the above-mentioned subdirectories: "http://www.companyA.com/products/Events/" and "http://www.companyA.com/products/Ads/." The list of URL patterns 510 in each site tracking list record 505 allows a remote site to enumerate specific directories, for example, in which they would like to track user hits.

The advantages of providing forms of URL patterns with wild cards becomes apparent with reference to the directory structure 600. Assume the "*" character is a wild card. That is, it matches zero or more characters. Since, CompanyA has two subdirectories with press releases, a convenient way to track hits in both is with the following URL pattern: "http://www.companyA.com/*press_releases/." Without the use of a wild card, the equivalent URL patterns are as follows: "http://www.companyA.com/press_releases/" and "http://www.companyA.com/products/press_releases." Thus, it should be appreciated that wild cards and regular expressions provide additional efficiency and convenience in the specification of URL patterns.

Hit Accumulation

FIG. 7 is a flow diagram illustrating a method of performing hit accumulation according to one embodiment of the present invention. In this embodiment, each site hit record 515 begins in an initial state having an indication of the remote site (e.g., the name 530) and a timestamp 525 representing the time at which hit accumulation began. Initially, the hit accumulation server 415 waits for an indication that a client request has been serviced from the proxy's local cache (step 710). For example, the hit accumulator processor 430 may determine that a new entry has been made to the hit log 420.

Upon receiving an indication that the proxy 400 has served up a cached response, the hit accumulation server 415 determines if the URL of the document retrieved from the proxy's local cache is one whose hits are to be tracked. As discussed above, not all hits are tracked. In this embodiment, hits are tracked only for documents matching URL patterns that have been registered in a tracking list such as the site tracking list 425, discussed above. Therefore, the hit accumulator processor 430 compares the URL of the retrieved document to URL patterns 510 in each site tracking list record 505 to determine if the hit will be recorded in the per site hit database 440 (step 720). If no URL patterns 510

match the retrieved document the hit is ignored. Otherwise, if the retrieved document matches any of the URL patterns 510, then the appropriate site hit record 515 in the per site hit database 440 is updated (step 730).

Update of the site hit record 515 can be explained briefly with respect to FIG. 5B. In this embodiment, the appropriate site hit record 515 is searched for an entry that matches the URL of the retrieved document. If the retrieved document's URL does not already exist in the list of hits 520 for the site hit record 515, then the URL is added and its count is set to one since this is the document's first hit. However, if the retrieved document's URL was already in the list of hits 520 (meaning it has had at least one previous hit), then only the corresponding count needs to be incremented. In this manner, each document retrieved from the proxy's local cache that matches a tracked URL pattern will have an entry in the list of hits 520 with a corresponding count indicating the number of cache hits.

Hit Reporting

Referring now to FIG. 8, a method of hit reporting according to one embodiment of the present invention is illustrated. In this embodiment of the present invention, the hit accumulator server 415, in addition to its other responsibilities, acts as a Web server providing an HTTP interface by which remote site administrators can access the accumulated hits for their respective tracked sites. The hit report processor 450 waits until a request is received from a remote site administrator (step 810). Preferably, the HTTP address on the hit accumulation server 415 can be used to identify the requester of the information. For example, the hit report for Company A, might be accessed on the hit accumulation server 415 at: "http://www.webtv.net/hits/company_a."

To limit access to the hit reports a secure communication technology such as Secure Sockets Layer (SSL) or other available secure communication protocol can be used to keep the hit information private by providing encrypted communications across the network. Additionally, the report requests can be authenticated to assure only a particular remote server or individual can access the information (step 820).

Once a request has been received from a remote site administrator and it has been optionally authenticated, then a report can be generated from the hit data accumulated such as the list of hits 520 for the particular site hit record 515 (step 830). In this embodiment, the report may include a list of URLs and their corresponding counts since the last report. For convenient access via the Web, the report may be formatted in an HTML format. Also, for the convenience of the remote site administrator a timestamp that identifies the starting point of the accumulation may be included in the report. The level of specificity of the URL list may be at the document level thereby allowing the remote site administrator to determine the number of hits for individual documents, however, it may also be helpful to additionally summarize the hits by directory, for example. It will be recognized that numerous other ways of formatting and arranging the hit reports are possible.

After the report has been formatted, the response containing the report is transmitted to the remote site administrator (step 840).

In this embodiment, before resuming the hit accumulation of FIG. 7, the accumulated data in the site hit record 515 is cleared (step 850) also the timestamp 525 is reset to reflect the current time. The above steps for retrieving a report from

the proxy may be periodically repeated at the convenience of the remote site administrator whenever an accurate total hit count is desired.

In alternative embodiments, hit reports may be provided to remote sites in a number of other ways. Hit reports need not be initiated by a request from the remote site administrator. For example, the proxy may periodically send unsolicited hit reports via e-mail, the proxy may periodically download hit updates to a device specified by the remote site administrator, or the hit reports might be transmitted to remote site administrators in the form of datagrams. In any event, the assignees of the present invention appreciate a variety of reporting mechanisms are possible.

Allocation of Cache Space Within a Proxy

FIG. 9 is a data flow diagram illustrating the interaction of proxy components according to another embodiment of the present invention. In this embodiment, proxy 900 includes a plurality of proxy servers 405 communicatively coupled to a dispatcher 910 and a hit accumulator server 415. Rather than allowing a given proxy server's cached contents to be determined based upon the requests of an associated client, the content of the Web can be distributed among proxy servers 405 by a hash algorithm executed by the dispatcher 910. The hash algorithm preferably maps a given URL to one and only one of the plurality of proxy servers 405. This can be accomplished using a portion of the output of a secure hash algorithm such as the Message Digest 5 (MD5) hash algorithm. The hash algorithm can be thought of as a mechanism for assigning a range of URLs to each of the proxy servers 405 in the proxy 900.

In this embodiment, the dispatcher 910 receives document requests including URLs from a client such as WebTV™ client 1. Based upon the URL in the request, the dispatcher determines the proxy server 405 in which the document should be cached and forwards the client request to that proxy server 405. If the document requested by the client is not found in the proxy server's local document cache 465, then the proxy server 405 requests the document from the appropriate server (e.g., a remote server) and caches the document when it is received from the server.

If redundancy is desired, the hashed result of a URL may be used to identify a cluster of two or more proxy servers rather than a single proxy server 405. In this manner, the load required to support a popular document can be shared among a group of proxy servers.

In an alternative embodiment, a decentralized dispatching scheme can be implemented. For example, the proxy servers 405 may be arranged to form an interconnected ring configuration and the functionality of the dispatcher 910 may be incorporated into each proxy server 405. In this embodiment, the client document requests may be initially handled by one of the proxy servers 405 in the ring. If the requested document is not found in the local cache of the initial proxy server, the initial proxy server may forward the request to the appropriate proxy server based on the hashing scheme discussed above.

FIG. 10 is a flow diagram illustrating a method of dispatching requests to segregate the storage of documents according to one embodiment of the present invention. While both a centralized and a decentralized request dispatching mechanism have been discussed above, the method described below is generally applicable to both. In this embodiment, initially, a document request is received from a client (step 1010).

If a centralized dispatcher such as dispatcher 910 receives the request, then based upon the URL an appropriate proxy

server is determined based upon the output of the hash algorithm (step 1020).

However, in a decentralized dispatching environment, the initial proxy server receiving the client request may assume it is the appropriate proxy server and first check its local document cache 465. If the document is not present, then proxy server may perform the hash algorithm on the URL to determine which of the remaining proxy servers is appropriate for the request (step 1020).

After determining the proxy server appropriate for the client request, the request is forwarded to that proxy server (step 1030). The proxy server 405 attempt to service the request from its local document cache 465. If a cache hit occurs, then the document is immediately available from the proxy server's local document cache 465. However, if a cache miss occurs, the proxy server 405 will retrieve the document from an appropriate server and store a copy locally. In any event, the centralized or decentralized dispatching mechanism ultimately receives a response from the server (e.g., the document requested by the client) (step 1040). Finally, the response, typically in the form of an HTML document is forwarded to the client (step 1050). This method of caching documents segregates the content of the Web based upon the URL of the documents. Since each URL will map to only one proxy server 405, advantageously this approach more efficiently allocates the proxy's cache space by avoiding unnecessary redundancy.

In the foregoing specification, the invention has been described with reference to specific embodiments thereof. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the invention. The specification and drawings are, accordingly, to be regarded in an illustrative rather than a restrictive sense.

What is claimed is:

1. A method of tracking hits in a proxy, the proxy including a document cache having stored therein recently requested documents, the proxy coupled to a client and to a remote server, the method comprising the steps of:

the proxy maintaining information regarding client requests that are serviced from the document cache;
the proxy receiving a client request for a document;
the proxy determining whether to service the client request from its document cache or whether to forward the client request to another server;
updating a count if the client request is serviced from the document cache; and
the proxy providing the information to a remote site administrator.

2. The method of claim 1, wherein the information includes a count representing the number of times client requests for a particular document have been serviced from the document cache.

3. The method of claim 2, wherein the information includes a timestamp identifying a time period to which the count corresponds.

4. The method of claim 1, wherein the remote site administrator requests the information from the proxy thereby initiating the step of providing the information to a remote site administrator.

5. The method of claim 4, wherein the information provided to the remote site administrator is in the form of an Hypertext Mark-up Language (HTML) report.

6. The method of claim 4, wherein the proxy authenticates the remote site administrator's request prior to the step of providing the information to a remote site administrator.

7. The method of claim 1, wherein the step of providing the information to a remote site administrator further includes the step of the proxy transmitting unsolicited information from the proxy to the remote site administrator.

8. In a system having a hit accumulation server and one or more proxy servers, each of said one or more proxy servers including a local cache having stored therein one or more cached documents, the proxy coupled to a client and to a remote server, a method of tracking requests for documents stored in a proxy comprising the steps of:

a proxy server receiving a client request for a document;
the proxy recording a hit for the document if the document is available from the proxy server's local cache, the proxy server notifying the hit accumulation server that the client request was serviced from the local cache, and the hit accumulation server recording the hit and a path of the document to a table if the document corresponds to one of a set of monitored Uniform Resource Locator (URL) patterns; and

providing an indication of the number of hits for the documents to a remote site administrator.

9. The method of claim 8, wherein the step of the proxy server notifying the accumulation server further includes the steps of:

the accumulation server monitoring a common storage device, the common storage device accessible to the one or more proxy servers; and

the proxy server logging an entry to the common storage device if the proxy server services the client request from its local cache, the entry including a URL for the document requested by the client.

10. The method of claim 9, wherein the set of monitored URL patterns represents one or more directories on a remote site for which document hits are to be tracked, the method further including the steps of:

the accumulation server detecting the entry; and
the accumulation server comparing the URL in the entry to the set of monitored URL patterns to determine whether or not to record the hit.

11. A method of tracking hits by a proxy, the proxy including one or more proxy servers and a hit accumulation server, each of said one or more proxy servers including a local cache having stored therein recently requested documents, the proxy coupled to a client and to a remote server, the method comprising the steps of:

a proxy server receiving a client request for a document;
the proxy server determining whether to service the client request from its local cache or whether to forward the client request to another server;

inserting a new log entry onto a common storage if the proxy server services the client request from its local cache, the new log entry including a location of the document;

the accumulation server detecting the new log entry by monitoring the common storage;

the accumulation server comparing the location of the document to a predetermined set of directories to determine whether or not to record the hit;

recording the hit if the location of the document matches a directory in the predetermined set of directories; and
providing the number of hits for a set of documents located in a first subset of directories of the predetermined set of directories to a remote site administrator.

12. The method of claim 11, wherein the common storage comprises a hit log.

13

13. The method of claim 11, wherein the location of the document comprises a Uniform Resource Locator (URL).

14. A machine-readable medium having stored thereon data representing sequences of instructions, said sequences of instructions which, when executed by a processor, cause said processor to perform the steps of:

maintaining information regarding client requests that are serviced from a document cache of a proxy server;
receiving a client request for a document;
determining whether to service the client request from the document cache or whether to forward the client request to another server;

14

updating a count if the client request is serviced from the document cache; and

providing the information to a remote site administrator.

15. The machine-readable medium of claim 14, wherein the information includes a count representing the number of times client requests for a particular document have been serviced from the document cache.

16. The machine-readable medium of claim 15, wherein the information includes a timestamp identifying a time period to which the count corresponds.

* * * * *

UNITED STATES PATENT AND TRADEMARK OFFICE
CERTIFICATE OF CORRECTION

PATENT NO. : 5,935,207
DATED : August 10, 1999
INVENTOR(S) : Jay D. Logue, Lee S. Mighdoll

It is certified that error appears in the above-identified patent and that said Letters Patent is hereby corrected as shown below:

Col. 5, line 28, before "the remote" delete "the site"

Col. 5, line 47, after "retrieved" change "form" to --from--

Col. 6, line 52, after "such as" insert --a--

Col. 6, line 64, after "tracking list" change "435" to --425-- (See figure 4)

Col. 7, line 41, after "three" change "sit" to --site--

Col. 7, line 60, after "the" change "ad1.html" to --ad1.html--

Col. 7, line 62, after "and" change "Q1.html" to --Q1.html--

Col. 9, line 47, after "particular" change "sit" to --site--

Col. 10, line 28, after "Digest" change "5 (MD5)" to --5 (MD5)--

Col. 11, line 6, after "then" insert --the--

Signed and Sealed this

Twenty-sixth Day of September, 2000

Attest:



Q. TODD DICKINSON

Attesting Officer

Director of Patents and Trademarks